arXiv:1607.00869v1 [cs.AI] 4 Jul 2016

Modeling of Item-Difficulty for Ontology-based Multiple-Choice Questions

Vinu E.V[1], Tahani Alsubait[2] and P Sreenivasa Kumar[1]

[1] Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India
[2] Information Science Department, Umm Alqura University, Makkah, Saudi Arabia
{vinuev, psk}@cse.iitm.ac.in,
tmsubait@uqu.edu.sa

Abstract. Multiple choice questions (MCQs) that can be generated from a domain ontology can significantly reduce human effort & time required for authoring & administering assessments in an e-Learning environment. Even though there are various methods for generating MCQs from ontologies, methods for determining the difficulty-levels of such MCQs are less explored. In this paper, we study various aspects and factors that are involved in determining the difficulty-score of an MCQ, and propose an ontology-based model for the prediction. This model characterizes the difficulty values associated with the stem and choice set of the MCQs, and describes a measure which combines both the scores. Furthermore, the notion of assigning difficultly-scores based on the skill level of the test taker is utilized for predicating difficulty-score of a stem. We studied the effectiveness of the predicted difficulty-scores with the help of a psychometric model from the Item Response Theory, by involving real-students and domain experts. Our results show that, the predicated difficulty-levels of the MCQs are having high correlation with their actual difficulty-levels.

Keywords: MCQ generation, ontologies, assessment tests, difficulty-level of MCQs

1 Introduction

Finding the difficulty-level of computer-generated questions, has recently gained much attention in educational community as well as in computer science community. This is mainly due to the wide-spread use of computer based tutoring systems in guiding and assessing learners. When it comes to assessing the knowledge level of a learner using an evaluation test — preliminary and concluding evaluations — the difficultly-levels of the questions used will play a crucial role in classifying the learners based on their knowledge proficiencies [1].

An assessment is mainly conducted to determine the amount and quality of knowledge a learner has gained from a session (or a series of sessions) of knowledge transfer. Therefore, one of the desired properties of an assessment test is to discriminate between good, average and low level learners [2,3]. Reducing burden in authoring and administering a test is another desired property of an assessment. Usually for large-scale assessments, multiple choice questions (MCQs) are of great usefulness, as the test administration and evaluation process can be automated using a computer program.

In the literature, there have been many attempts such as [4,5,6,7] for generating MCQs. These approaches have used structured or unstructured knowledge sources for

generating MCQs. Recent works were focusing more on structured knowledge sources like Web Ontology Language (OWL) ontologies, than unstructured knowledge sources (e.g., text file), for many reasons. Some of the reasons for this trend that are suggested by Alsubait et al. [8] are the richness of ontologies and their resemblance to knowledge representations acquired by learners in the course of their studies. In addition, OWL ontologies are based on a decidable fragment of first order logic, called Description Logics [9], where reasoning can be done to infer implicit relations among the domain entities. This inferencing process is similar to the reasoning carried out by a learner to reach conclusions on a subject matter.

Even though there are numerous methods for generating MCQs, methods to estimate difficulty-levels of such MCQs are less explored. Without a proper difficulty-level estimation method, the generated MCQs are less employable in pedagogy applications.

In this paper, we propose an automated method to determine the difficulty-levels of MCQs that are automatically generated from ontologies. Furthermore, we discuss how the knowledge proficiency of a learner affects the difficulty-level of a question. Section 5 describes a method to check the sanity of the predicted difficulty values of the MCQs. In Section 7, we describe an experiment which is based on a psychometric model, to determine the actual difficulty-levels of a set of machine-generated MCQs. We then correlate this actual difficulty-levels with their estimated (or predicated) difficulty-levels, to examine the effectiveness of our approach.

2 Background and Motivation

2.1 Ontology-based MCQ generation systems

Ontology-based MCQ generation was first introduced by Papasalouros et al. [10], where they have introduced a set of strategies based on different ontological entities for framing MCQs and the corresponding distracting answers. Their method lacks proper theoretical support for when to use which strategy, and the stem of all the generated questions remains the same (*"Choose the correct sentence"*). Later, Tosic and Cubric [4] generated MCQs of knowledge level (*"Which of the following definition describes the concept* C*?"*), comprehension level (*'Which one of the following response pairs relates in the same way as* a *and* b *in the relation* R*?'*), application level (*"Which one of the following examples demonstrates the concept* C*?"*) and analysis level (*"Analyze the text* x *and decide which one of the following words is a correct replacement for the blank space in* x*."*). They further extended the work introduced by Holohan et al. [11], by introducing stems that use annotation information in the ontology. Strategies similar to Papasalouros's strategies were adopted to find the distractors for the generated question statements. Another MCQ generation approach was by Alsubait et al. [12]. Their approach was limited to analogy type questions. Later, they extended their work to include different kinds of MCQs [13,14]. Other than the above MCQ generation approaches, there are works such as [15,16], which utilize simple ontology statements: concept inclusions, role hierarchy and (concept and role) assertions, to generate basic domain related questions.

Most recent works on MCQ generation can be found at [5,7,14]. In [5], the authors make use of the semantics of the domain, represented in the form of TBox axioms along with ABox axioms, to frame interesting MCQs. Vinu and Kumar [7] illustrated an effective and practically adoptable generic (i.e., domain independent) method to generate MCQs from the assertional facts of a given ontology. Their method, *pattern-based*

MCQ generation, utilizes different combinations of predicates (i.e., concept names and roles) associated with the instances of an ontology, for framing the stems. As at present, the pattern-based approach is the state-of-the-art technique and since we can generalize existing methods to this approach, we predict the difficulty values of the MCQs that are generated using that method. A detailed explanation of the pattern-based MCQ generation method is given in the next subsection to set the context for the current work.

2.2 Pattern-based MCQ generation

In the pattern-based MCQ generation, introduced in [7], components of an MCQ — *Stem* and *Choice set* (i.e., key and distractors) — are generated using simple SPARQL templates. In this approach, a stem can be considered as a set of *conditions* which asks for a solution which is explicitly present in the ontology. The set of conditions in a stem is formed using different combinations of (unary or binary) predicates associated with an instance in an ontology. Example-1 is such an MCQ, which is framed from the following assertions that are associated with the (*key*) instance *birdman*. For ease of reading, all examples presented in this paper are from the Movie domain.

Movie(birdman) //concept assertion
isDirectedBy(birdman,alejandro) //role assertion
hasReleaseDate(birdman,"Aug 27 2014") //role assertion

Example 1. **Stem:** Choose a *Movie*, which *is directed by Alejandro* and *has release date Aug 27, 2014*.

Choice set:	a. Birdman	b. Titanic	c. Argo

The possible predicate combinations (or question patterns) of size[3] one w.r.t. an instance x can be denoted as: $x\ \overrightarrow{O}\ i$, $x\ \overleftarrow{O}\ i$, $x\ \overrightarrow{D}\ v$ and $x\ \overrightarrow{d}\ C$, where i is an instance, $\overrightarrow{d}$ is *rdf:type*, $\overrightarrow{O}$ and $\overleftarrow{O}$ represents object properties that are having different directions, $\overrightarrow{D}$ denotes datatype property, v stands for the value of the datatype property and C is a class name. We call the instance x as the *pivot-instance* of the question pattern. The arrows ($\leftarrow$ and $\rightarrow$) represent the directions of the properties w.r.t. the pivot-instance. In [7], the authors studied the pattern combinations of sizes one and two; but, in this paper we are generalizing the pattern size to n (a positive integer). A matching tuple of the pattern (of size two): $i_1\ \overrightarrow{O_1}\ x\ \overleftarrow{O_2}\ i_2$, is represented as a list: $[(a, O_1, b), (c, O_2, b)]$, where b corresponds to the pivot-instance x. Linguistic representations of tuple lists are addressed as the *stems*. For a pattern of size n, there will be n ordered triples in the list. These triples without the pivot-instance is considered as the conditions of a stem. The conditions of a stem S is represented as $C_S = \{r_1, r_2, ..., r_n\}$, where, for an instance y, "$r_i.y$ is true" means that y satisfies the condition r_i. In this paper, we consider the pivot-instance of a tuple, as the *key* corresponding to the stem generated from that tuple. We do not consider tuples containing blank nodes for stem generation.

The distractors for these MCQs are selected from the set of instances (or literals) of the ontology which belong to the intersection classes of the domain or range of the predicates that are present in the stem. In our earlier work [7], this set was addressed as

[3] Signifies the number of predicates in a combination

Potential-set. For example, the potential-set corresponding to the pattern $i_1 \overrightarrow{O_1}\, x \,\overleftarrow{O_2}\, i_2$, is *Range*$(O_1) \sqcap$ *Range*(O_2). Readers are referred to the original paper, for more details.

The *choice set* of an MCQ can be considered as a set of instances from its potential-set along with the key, presented in some random order. In this work, we consider the cardinality of the choice set as three (including exactly one key), as it is the optimal number of options preferred for conducting an MCQ test by many researchers [17,18].

2.3 Difficulty-level calculation

A simple notion to find the difficulty-level of an ontology-generated MCQs was first introduced by Cubric and Tosic [4]. Later, in [8], Alsubait et al. extended the idea and proposed a similarity-based theory for controlling the difficulty of ontology generated MCQs. In [12], they have applied the theory on analogy type MCQs. In [13], the authors have experimentally verified their approach in a student-course setup. The practical solution which they suggested to find out the difficulty-level of an MCQ is w.r.t. the degree of similarity of the distractors to the key. If the distractors are very similar to the key, students may find it very difficult to answer the question, and hence it can be concluded that the MCQ is difficult. For instance, the following table shows the similarity of the distractors with the key, for a given stem (see Example-1 for stem details), where the difficulty due to distractors = $(0.5 + 0.4)/2 = 0.45$.

Choice set:	a. Birdman (*Key*)	b. Titanic	c. Argo
Similarity with the key:	1.0	0.5	0.4

Similarity-based theory [8] was the only effort which is utilized in the literature [5,7,13], to calculate or control the hardness of an ontology-based generated MCQ. The difficulty value determined by similarity-based theory considers only the similarity of the correct-option (*key*) with the wrong-options (*distractors*). That is, if the similarity is high, then it implies that the MCQ has a high difficulty-level and vice versa.

In many a case, the stem (question statement) of an MCQ is also a deciding factor for the hardness of an MCQ. For instance, the predicate combination (discussed in the previous subsection) which is used to generate a stem can be chosen such that they can make the MCQ harder or easier to answer. We investigate this aspect in Section 3. Difficulty-level calculation based on Similarity-based theory is detailed in Section 5. A unified method, by combining the two methods is proposed in Section 6. An empirical study in a classroom setup (see Section 7) was done to evaluate the effectiveness of the proposed method.

3 Difficulty-level of Stem

Consider the stems (underlined portions indicate the concept names and roles used):

Stem-1: *Choose a Movie which was directed by Clint Eastwood*.
Stem-2: *Choose an Oscar movie which was directed by Clint Eastwood*.

In the empirical study that we have conducted as a part of our earlier work, we observed that, the former question, for a given set of options (we will see the significance of the *choice set* later), is not correctly answered by many people with good knowledge about the movie domain, when compared to the latter question (especially when the choice set does not contain any famous movie names). A rational explanation for this behavior is that, *Explanation-1:* in the latter question, it is given that the movie is an Oscar Movie, which is a *good hint* for an expert to correctly answer the question, whereas in the former case, she will find it difficult to choose the right answer, as the hint is very poor in identifying the key. It should be noted that, this explanation is given in reference to the perspective of an expert-level learner.

On the other hand, if we look from a beginner-level learner's perspective, she may identify the hint; but may not be able to utilize it due to lack of expertise. For her, the difficulty of a question will increase if the stem contains specific details of the domain. If we revisit the previous stems, Stem-2 is likely to be more difficult than Stem-1; this cannot be established using Explanation-1. Therefore, it is necessary to come up with a new explanation for question difficulty which should go with the beginner-level perspective.

A logical explanation in reference to a beginner-level learner can be, *Explanation-2:* when compared to Stem-1, Stem-2 requires the additional knowledge of Oscar movies for correctly answering the question. Therefore, the latter question is more difficult for a beginner-level learner than the former one. But, here we are not claiming that Stem-1 is easy for beginner-level learners; rather, we are saying that Stem-2 is relatively harder to answer than Stem-1, fulfilling the following desired aspect.

A *desired aspect* of a difficulty-level of a question is that, if a given question is difficult for an expert-level learner then, it should also be difficult for a beginner-level learner.

A proposition that can be made out from the aforementioned explanations based on the perspective of a learners knowledge-level is that, *difficulty level of an MCQ is relative to the knowledge proficiency of a test taker*. In other words, the difficulty-score (value) of an MCQ should be defined in terms of the knowledge proficiency of the test takers. This indeed makes it necessary to have multiple difficulty-scores for a given MCQ. In this paper, we limit the scope of our studies by categorizing the knowledge proficiencies (a.k.a. *trait-levels*) of learners into three: *beginner*, *average* and *expert*. We then formulate difficultly measures for each of the three cases.

3.1 Difficulty measure for expert-level learner

As we have mentioned above, the difficulty-score of an MCQ w.r.t. to an expert learner depends on how good the hints are helping her to answer the MCQ correctly. By hints we mean the conditions present in the stem. The obvious method to measure the quality of a hint is by measuring how well the *answer space* of the question is affected by the hint; in the current context, answer space refers to the number of instances that satisfy a given condition. That is, if the answer space of the question is greatly reduced because of a particular hint, then it is a *good hint*. For example, in the stems that we saw earlier, *directed by Clint Eastwood* is a good hint — as the count of instances which satisfy this condition is low. Since both stems have the same (good) hint, we are not gaining any insight about their relative hardness, just by looking at these good hints alone. We have to look at the other predicates (here, the unary predicates: *Movie* and *Oscar Movie*)

as well. These predicates clearly distinguish the two stems (i.e., Stem-2 is easier to answer than Stem-1), by looking at the count of instances satisfying these predicates. This observation, made us to conclude that, *answer spaces of all the predicates in a stem should be considered for finding its difficulty-score.*

Now, consider another stem, **Stem-3:** *Choose a Thriller Movie which was directed by Clint Eastwood*. If the answer spaces for the predicates: *Oscar Movie* and *Thriller Movie*, happened to be of same cardinality, then, both stems (Stem-2 and Stem-3) will have the same difficulty-scores. On the contrary, we observed that, Stem-2 is less difficult for an expert to answer correctly than Stem-3. A rationale for this behavior is that, the concept *Oscar Movie* is a more popular concept than *Thriller-Movie*. Therefore, we consider, *the popularities of the predicates present in a stem as another factor which affects its difficulty.*

We propose the difficulty-score of a stem S, which is having predicates (concepts or roles) $P = \{p_1, p_2, ..., p_n\}$ (a.k.a *Property set*), generated from an ontology $\mathcal{O}$, in terms of the answer spaces (denoted as $ASpace_{\mathcal{O}}(.)$), and popularities of these properties (represented as $PredicatePopularity_{\mathcal{O}}(.)$), as shown below:

$$D_{expert}(S, \mathcal{O}) = \sum_{i=1}^{n} \frac{\log(1 + |ASpace_{\mathcal{O}}(p_i, S)|)}{PredicatePopularity_{\mathcal{O}}(p_i, S)} \quad (1)$$

Answer space. The answer space of p in S (i.e., $ASpace_{\mathcal{O}}(p, S)$) is the set of instances which satisfy the predicate p in $\mathcal{O}$. If p is a unary predicate (i.e., a concept), then the definition is straight forward. But, to find the answer space of a binary predicate, we consider the directionality of the predicate w.r.t. the pivot-instance of the corresponding pattern and accordingly choose the subject or object along with the predicate (this is called the *concept equivalent* of the binary predicate). For example, consider the tuple corresponding to Stem-3: [(*mystic_river, is_a, Thriller_movie*), (*mystic_river, isDirectedBy, clint_eastwood*)], we find the instances which belong to the concept *Thriller_movie* and those instances which satisfy the constraint (i.e., the concept equivalent): $\exists$*isDirectedBy.* $\{$*clint_eastwood*$\}$, to find the answer space. For a tuple of the form: [(*mystic_river, is_a, Thriller_movie*), (*clint_eastwood, isDirectorOf, mystic_river*)], since the directionality of the predicate in the second triple is different from that in the tuple for Stem-3, we take the concept equivalent as: $\exists$*Inv(isDirectorOf).*$\{$*clint_eastwood*$\}$ — if *Inv(.)*, inverse of the predicate is not present, we look for the pattern (using the SPARQL query fragment:) "*clint_eastwood isDirectorOf ?x*" in the rdf graph.

Predicate Popularity. We calculate the popularity of a predicate p in S (i.e., $Predicate\ Popularity_{\mathcal{O}}(p, S)$) by finding the mean of the popularities of all the instances which satisfy p in $\mathcal{O}$; the popularity of an instance i, satisfying p is denoted as $Popularity_{\mathcal{O}}(p, i)$. Here, p can be a unary predicate or a binary predicate. If p is a binary predicate, its concept equivalent will be considered. A widely used measure for finding the popularity of a concept c is based on the count of the instances of the other concepts that are related to the instances that satisfy c [19]. We make use of this measure to find the popularity of a concept. Therefore, $PredicatePopularity_{\mathcal{O}}(.)$ of a unary predicate p of a stem S can be calculated as follows, where n is the cardinality of the set $ASpace_{\mathcal{O}}(p, S) = \{i_1, i_2, ..., i_n\}$:

$$PredicatePopularity_{\mathcal{O}}(p, S) = (1/n) \times \sum_{j=1}^{n} Popularity_{\mathcal{O}}(p, i_j) \quad (2)$$

In Eq. 2, $Popularity_{\mathcal{O}}(p, i)$ gives the count of the instances that are related to i by a relation R, such that they do not satisfy p. That is, $Popularity_{\mathcal{O}}(p, i) = |\{j|\mathcal{O} \models p(i) \sqcap R(j, i) \wedge \mathcal{O} \not\models p(j)\}|$.

3.2 Difficulty measure for beginner-level learner

For a beginner-level learner, the difficulty-level depends on how detailed the question is. Intuitively, if the stem contains domain specific conditions, the probability of a learner for correctly answering the question will reduce. Therefore, we relate the depths of the concepts and roles that are used in the stem to the concept and role hierarchies of the ontology, to determine the stem difficulty. To achieve this, we introduce $DepthRatio$ for each predicate p in an ontology. $DepthRatio$ is defined as:

$$DepthRatio_{\mathcal{O}}(p) = \frac{\text{Depth (or length) of } p \text{ from the root of the hierarchy}}{\text{Maximum length of the path containing } p} \tag{3}$$

For a stem S, generated from an ontology $\mathcal{O}$, with x as its pivot-instance and P as its property set, let $\mathcal{C}$ denote the set of concepts satisfied by x, and let $\mathcal{R}$ represents the set of roles such that either x is present at their domain (subject) or range (object) position $\big($i.e., $R \in \mathcal{R} \implies \mathcal{O} \models R(x, i) \vee R(i, x)$, where i is an arbitrary instance in $\mathcal{O}\big)$. For each $p \in P$, we find the largest subset in $\mathcal{C}$ (if p is a concept) or we find the largest subset in $\mathcal{R}$ (if p is a role), such that the elements in the subset can be related using the relation $\sqsubseteq$, and p is an element in that subset. The cardinality of such a subset forms the denominator of Eq. 3, and the numerator is the position of the predicate p from right (right represents the top concept or top role), when the elements in the subset are arranged using the relation $\sqsubseteq$.

In addition to the $DepthRatio$, answer space of the predicates in a stem also affect the difficulty-score of an MCQ (recall Stem-1 and Stem-2 in Section 3). That is, if the cardinality of the answer space increases, there is a tendency that stem becomes more generic. For example, if we include the condition $\exists$*isDirectedBy.*$\{$*tim_miller*$\}$ in the stem rather than $\exists$*isDirectedBy.*$\{$*clint_eastwood*$\}$, the answer space will largely reduce (as Tim Miller is a young director who has directed only a few movies), resulting in an increase in the difficulty-score. Therefore, we define the difficulty-score of a stem having n number of predicates, in terms of the depth ratios and answer space of its predicates, as follows.

$$D_{beginner}(S, \mathcal{O}) = \sum_{i=1}^{n} DepthRatio_{\mathcal{O}}(p_i) \times \frac{1}{1 + \log(1 + |ASpace_{\mathcal{O}}(p_i, S)|} \tag{4}$$

The portion of the equation: $1/(1 + \log(1 + |ASpace_{\mathcal{O}}(p_i, S)|))$, shows that the difficulty-score of a stem is inversely proportional to the cardinalities of the answer spaces of its predicates. Including a factor of depth ratio in Eq. 4, will help to make the difficulty-score high, when the stem contains more specific predicates.

3.3 Difficulty measure for average-level learner

For a learner with average-level knowledge proficiency, both the factors which we discussed in Section 3.1 and Section 3.2, need to be considered for calculating the stem difficulty. Therefore, we take the mean of $D_{beginner}$ and D_{expert}, to find $D_{average}$.

$$D_{average}(S, \mathcal{O}) = (1/2) \times \big(D_{expert}(S, \mathcal{O}) + D_{beginner}(S, \mathcal{O})\big) \tag{5}$$

4 Validity of stems

As we have seen above, given an MCQ stem S, generated from an ontology $\mathcal{O}$, we can have three difficulty-scores (i.e., $D_{expert}(S, \mathcal{O}), D_{average}(S, \mathcal{O}), D_{beginner}(S, \mathcal{O})$) assigned to it, corresponding to expert, average and beginner level learners, respectively. Let d_e, d_a and d_b be the normalized[4]scores ($\rightarrow [0, 1]$) of $D_{expert}, D_{average}$ and $D_{beginner}$, respectively. Since d_a is the average of d_e and d_b, we can represent the difficulty-score of S in terms of the tuple (d_e, d_b). Based on the notion that if an expert finds a stem to be difficult then it is difficult for a beginner as well, and based on the desired property of a good assessment test (i.e., it should be able to distinguish expert, average and beginner level learners), we fix the validity definition of a given MCQ stem as:

Definition 1. *Validity of a stem. A given stem S, with difficulty-score tuple (d_e, d_b), is* valid*, if and only if $d_e < d_b$.*

We ensure that the generated MCQs are valid as per the above definition before they are used in the experiment for empirical evaluation.

5 Difficulty-level due to Choice set

5.1 Similarity as a factor

In the similarity based method (see Section 2.3), the similarity measure for finding the instances' degree of similarity is subjective to the employed stem generation technique. For instance, in our earlier work [5], we proposed the *Label-set similarity ratio* for finding the closeness of instances (where we associate TBox based restrictions that are present in the stem, to the instances to calculate their similarity); this measure cannot be used for those MCQs which are generated using pattern-based methods, since TBox axioms are not directly involved in the stem generation.

Now the question is, how to calculate the similarity of two instances w.r.t. a stem S from a given ontology $\mathcal{O}$. Since there are no measures in the literature which satisfy our requirement, we propose a similarity measure, called *Instance Similarity Ratio*, in Def. 2.

Given an ontology $\mathcal{O}$, two instances k (the key) and d (a distractor), and a stem S, to calculate the Instance Similarity Ratio, we consider (1) similarity of the instances w.r.t. the conditions (represented as $C_S = \{r_1, r_2, .., r_n\}$) in S. (The conditions in S, that are satisfied by k and d, can be represented as $C_k = \{r | r.k$ is true in $\mathcal{O} \wedge r \in C_S\}$ and $C_d = \{r | r.d$ is true in $\mathcal{O} \wedge r \in C_S\}$ respectively, where $r.x$ is true if the condition r is satisfied by the instance x.); (2) generic similarity of the two instances w.r.t. all the conditions they satisfy. (The set of all the conditions satisfied by an instance i in $\mathcal{O}$ is represented as $A_i = \{r | r.i$ is true in $\mathcal{O}\}$.)

Definition 2. *Instance Similarity Ratio, $Sim_{\mathcal{O}}(.)$, is defined as:*

$$Sim_{\mathcal{O}}(k, d) = \frac{\#\textit{Instances in } \mathcal{O} \textit{ that satisfy } C_k \cup C_d}{\#\textit{Instances in } \mathcal{O} \textit{ that satisfy } C_k \cap C_d} + GSim_{\mathcal{O}}(k, d)^2 \quad (6)$$

[4] Normalization is done by dividing the scores by the maximum score.

The first part of the right hand side denotes the similarity of k and d, w.r.t. the conditions in S. The second part, $GSim_{\mathcal{O}}(.)$, denotes the generic similarity of the instances.

$$GSim_{\mathcal{O}}(k,d) = \frac{1}{2} \times \frac{\#\textit{Instances in } \mathcal{O} \textit{ that satisfy } A_k + \#\textit{Instances in } \mathcal{O} \textit{ that satisfy } A_d}{\#\textit{Instances in } \mathcal{O} \textit{ that satisfy} A_k \cap A_d}$$

A_k and A_d represent the set of all the conditions that are satisfied by k and d respectively in $\mathcal{O}$.

The first part of Eq. 6 gives the ratio of the number of instances in $\mathcal{O}$ that satisfy all the conditions that are satisfied by k and d, to the number of instances that satisfy the conditions that are satisfied by both k and d. This formula is similar to the *Jacquard similarity* measure. To ensure that the selected distractor d has a minimum similarity with k, we consider $C_k \cap C_d \neq \phi$. The second part of the equation gives the square of the similarity of the two instances which is obtained by considering all the conditions they satisfy.

Definition 3. *Difficulty-score due to Choice set. The difficulty-score due to the choice set $X = \{k, d_1, d_2, d_3, ..., d_l\}$ (with k as the key) of an MCQ is defined as:*

$$D_{similarity}(X, \mathcal{O}) = (1/l) \times \sum_{i=1}^{l} Sim_{\mathcal{O}}(k, d_i) \tag{7}$$

5.2 Popularity as a factor

It is observed that the popularities of the instances that are being chosen as the key and as the distractors have a negative impact on the overall difficulty-score of the MCQ. For instance, consider the choice sets in Example 2. The key and the distractors in the first choice set have relatively high popularities (due to more number of connectivity) than those in the second choice set. Intuitively, it is more likely that the most popular options are easily be identified by a learner as a correct answer or as a wrong answer. To capture this notion, we came up with an equation to find the difficulty-score due to popularity of a choice set as follows:

Definition 4. *Difficulty-score due to popularity. The difficulty-score due to the popularities of the instances in a choice set $X = \{k, d_1, d_2, ..., d_l\}$ (generated from an ontology $\mathcal{O}$) where k is the key and $\{d_1, d_2, ..., d_l\}$ are the distractors, is defined as:*

$$D_{popularity}(X, \mathcal{O}) = \frac{1}{Popularity(X, \mathcal{O})} \tag{8}$$

The popularity of an instance x in X can be said as the **connectivity** *of x from the other individuals in $\mathcal{O}$ which belong to a concept which x does not belong to. $\mathcal{C}_{\mathcal{O}}(j)$ represents the connectivity (formally defined below) of an instance j in $\mathcal{O}$.*

$$Popularity(X, \mathcal{O}) = (1/2) \times Popularity_{\mathcal{O}}(k) + \sum_{i=1}^{l} \log\left(1 + \mathcal{C}_{t,\mathcal{O}}(d_i)\right) \tag{9}$$

In Eq. 9, the popularity is calculated by finding the sum of the connectivities of its instances, by giving more preference to the connectivity value of the key instance. — i.e., the sum of half the connectivity value of the k and $\log$ of the connectivity values of the distractors. This is because, given a popular instance as a key and less popular instances as distractors, a learner can easily figure out the correct answer, where as given a less popular instance as key and including a popular instance as distractor, still can deviate the learner from answering the MCQ correctly. For this reason, while calculating the difficulty due to distractors, we divide it with the overall popularities of the instances in the choice set, such that if the popularity of the choice set increases the MCQ will become more easier to answer.

Definition 5. *Connectivity. The connectivity of an instance x in ontology $\mathcal{O}$ is defined as follows, where C_x and C_y are concepts in $\mathcal{O}$.*

$$\mathcal{C}_{t,\mathcal{O}}(x) = \#\{y \mid \mathcal{O} \models C_x(x) \sqcap C_y(y) \sqcap R(y,x) \wedge \mathcal{O} \not\models C_y(x) \wedge \\ not(C_x \sqsubseteq C_y) \wedge not(C_y \sqsubset C_x)\}$$

The equation gives the count of the instances which are related to x by a relation R and whose satisfying concepts do not satisfy x and are not hierarchically (sub-class–super-class relationship) related to any of the satisfying concepts of x.

Example 2. **Stem:** Choose an *American martial arts film.*

Choice set 1:	a. The Karate Kid (*Key*)	b. 102 Dalmatians	c. Gone with the Wind
Popularity:	11	8	10

Choice set 2:	a. Sunland Heat (*Key*)	b. A Vanished World	c. Just Call Me Nobody
Popularity:	4	3	2

5.3 Overall difficulty-score due to choice set

The overall difficulty-score due to the choice set (denoted as DC) of an MCQ generated from an ontology $\mathcal{O}$ with stem S and choice set X is defined as:

$$DC(X, \mathcal{O}) = D_{similarity}(X, \mathcal{O}) \times D_{popularity}(X, \mathcal{O}) \tag{10}$$

6 Combined Score: Difficulty-level of MCQ

Having the difficulty-score of a stem S (denoted as DS) and the difficulty-score due to its choice set X (represented as DC), we can find the overall difficulty-score of the MCQ $Q = (S, X)$ (generated from ontology $\mathcal{O}$) by simply multiplying DS and DC. We choose to multiple the two scores, since, either a poor set of options or a easy stem can make an MCQ less difficult to answer.

$$D_{predicted}(Q) = DS(S, \mathcal{O}) \times DC(X, \mathcal{O}) \tag{11}$$

In the Eq. 11, depending on the pedagogical goal of the assessment (for example, whether the test MCQs are meant for identifying expert learners) we will choose $D_{beginner}, D_{average}$ or D_{expert}, as DS.

7 Empirical Evaluation

We have implemented a prototype system (developed using Java Runtime Environment JRE V1.6 and the Jena API[5] 2.11.0) which can generate MCQs and predict their difficulty-scores. The system is being tested by giving various ontologies (mentioned in [5,7]) as inputs. In this section, we mainly investigate the effectiveness of the proposed difficulty-score predication measures. For this purpose, we have generated test MCQs from a handcrafted ontology, and determined their difficulty-scores (a.k.a *predicted difficulty-scores*), using the method proposed in Section 6. Then, we compared these predicted difficulty-levels with their *actual difficulty-levels* that were estimated using principles in the Item Response Theory.

7.1 Estimation of actual difficulty-level

Item Response Theory is an item oriented theory which specifies the relationship between learners' performance on test items and their ability which is measured by those items [1]. In IRT, *item analysis* is a popular procedure which tells if an MCQ is too easy or too hard, and how well it discriminates students of different knowledge proficiencies. Here, we have used item analysis to find the actual difficulty-levels of the MCQs.

We make use of the simplest IRT model (often called *Rasch model* or the *one-parameter logistic model* (1PL)) for finding the actual difficulty-level of the MCQs. According to this model, we can predict the probability (P) of answering a particular item (with difficulty value α) correctly by a learner of knowledge proficiency level θ (a.k.a *trait level*), as specified in the following formula.

$$P(\theta, \alpha) = \frac{e^{(\theta-\alpha)}}{1 + e^{(\theta-\alpha)}} \tag{12}$$

Due to page limitation, we are providing the detailed theoretical background of the 1PL model in an online appendix[6]. To find the (actual) difficulty value, we can rewrite Eq.12 as follows:

$$\alpha = \theta - log_e\Big(\frac{P(\theta, \alpha)}{1 - P(\theta, \alpha)}\Big) \tag{13}$$

For experimental purposes, suitable θ values can be assigned for high, medium and low trait levels. Given the probability (of answering an MCQ correctly by learners) of a particular trait level, if the calculated α value is (approximately) equal or greater than the θ value, we can assign the trait level as its *actual difficulty-level*.

7.2 Experimental setup

A controlled set of MCQs from the DSA ontology has been used to obtain evaluation data related to its quality.

[5] https://jena.apache.org/

[6] https://www.overleaf.com/read/vygnvzgcjvrg (Date of access:28/04/2016)

DSA ontology. Data-structures & Algorithms (DSA) ontology (available at our project website[7]) models the various aspects of data-structures and algorithms. An initial version of this ontology was developed by Revuri and Kumar [20] for information retrieval application. Later, an enriched version of the ontology which is primely focusing on question generation application is prepared by our research team. The ontology is composed of 90 concepts, 30 object properties, 13 datatype properties and 165 individuals.

Question-set generation: For the experiment mentioned in this paper, we considered only three patterns (Pattern-1: $i\ \overrightarrow{O}\ x\ \overrightarrow{D}\ v$; Pattern-2: $i_1\ \overrightarrow{O_1}\ x\ \overrightarrow{O_1}\ i_2$; Pattern-3: $x\ \overrightarrow{O}\ i$ — as they are the common patterns) for generating MCQ stems from the DSA ontology. Using Pattern-1, we have generated 38 stems, only 16 stems were determined as valid stems. 368 stems were generated using Pattern-2, only 169 stems were predicated as valid stems. For Pattern-3, out of 405 generated stems, 185 were valid. These valid stems were then associated with choice sets such that their overall difficultly-level fall into the category of high, medium or low.

Test MCQs and instructions. We have administered a question-set of 24 (valid) test MCQs to 54 participants, with the help of a web interface. These 24 MCQs were chosen such that, the test contains 8 stems of each (predicated) difficulty-level (high, medium and low). The difficulty-scores of their stems were pre-determined using the method detailed in Section 6. Difficulty-levels (predicted difficulty-levels) were then assigned by statistically finding suitable range of values (corresponding to low, medium and high) from the obtained normalized difficulty-scores of all the stems. Suitable choice sets were then assigned to further tune their difficulty-levels. More details about the question-set preparation and the question-set itself are available at our project website. All the test MCQs were carefully vetted by human-editors to correct grammatical and punctuation errors, and to capitalize the proper nouns in question stems. Each MCQ contains choice set of cardinality three (with exactly one key) and two additional options: SKIP and INVALID. A sample MCQ is shown in Example-3.

Example 3. Choose an Internal Sorting Algorithm with worse case time complexity n exp 2.

Choice set:	a. Heap Sort	b. Bubble Sort	c. Breadth First Search
	d. SKIP	e. INVALID	

The responses from (carefully chosen) 54 participants — 18 participants of each trait level (high, medium and low) — were considered for generating the statistics about the item quality. The following instructions were given to the participants before starting the test.

1. The test should be finished in 40 minutes and all questions are mandatory.
2. You may tick the option "SKIP" if you are not sure about the answer. Kindly avoid guess work.
3. If you find a question invalid, you may mark the option "INVALID".
4. Avoid use of the web or other resources for finding the answers.
5. In the end of the test, you are requested to enter your expert level in the subject w.r.t this test questions, in a scale of high, medium or low. Also, kindly enter your grade which you received for the ADSA course offered by the Institute.

[7] https://sites.google.com/site/ontomcqs/item-difficulty-results (Date of access:28/04/2016)

Participant selection. Fifty four learners of the required knowledge proficiencies were selected from a large number of graduate level students (of IIT Madras[8]), who have participated in the online MCQ test. To determine their trait levels, we have instructed them to self assess their knowledge-confidence level on a scale of high, medium or low, at the end of the test. To avoid the possible errors that may occur during the self assessment of trait levels, participants with high and medium trait levels were selected from only those students who have successfully finished the course: CS5800: *Advanced Data Structures and Algorithms*, offered at the computer science department of IIT Madras. The participants with high trait level were selected from those students with either of the first two grade points (i.e., 10 - Excellent and 9 - Very Good). The participants with medium trait level were from those students who were having any of the next two grade points (i.e., 8 - Good and 7 - Satisfactory Work). The evaluation data collected for the item analysis is shown in Table 2.

Item analysis The probabilities of correctly answering the test MCQs (represented as P) by the learners are listed in Table 2. In the table, the learner sets L_1, L_2 and L_3 correspond to the learners l_1 to l_{18}, l_{19} to l_{36} and l_{37} to l_{54}, respectively. The learners in these learner sets have high, medium and low trait levels, respectively. The probability P (a.k.a., P_{qr}) of correctly answering an MCQ q for each of the learner sets (L_r) are obtained by dividing the count of learners in L_r who have correctly answered q by the cardinality of L_r. While calculating the P values, if a learner has chosen the option "SKIP" as the answer, the MCQ is considered as wrongly answered by her. If she has chosen "INVALID", we do not consider her poll for calculating P.

In Table 2, we list the α_i (actual difficulty) values that we have calculated using the corresponding P values. Eq.13 is used for finding the *difficulty* values. These α_i values were then used to assign the actual difficulty-level for the MCQs.

Table 1: Thumb rules for assigning difficulty-level, based on IRT

Trait level	α_i	Difficulty-level
High	(> 1.5) or $(\approx 1.5 \pm .45)$	High
Medium	(> 0) or $(\approx 0 \pm .45)$	Medium
Low	(> -1.5) or $(\approx -1.5 \pm .45)$	Low

We are particularly interested in the highlighted rows in Table 2, where an MCQ can be assigned an actual difficulty-level as shown in Table 1. That is, for instance, if the trait level of a learner is high and α_i is approximately equal to θ_l (ideally, $\alpha_i \geq \theta_l$), then, according to the IRT model, a difficulty-level of "high" can be assigned. In our experiments, to calculate α_i values for high, medium and low trait levels, we used θ_l values $1.5, 0$ and -1.5 respectively.

Results Out of the eight test MCQs which were having high predicted difficulty-levels, seven MCQs have high actual difficulty-levels. Among the eight medium difficulty-level MCQs, only one MCQ has deviated from its predicted difficulty-level. Three out

[8] https://www.iitm.ac.in/

of eight MCQs having low predicted difficulty-levels showed no correlation with their actual difficulty-levels. The overall correlation of the predicated difficulty-levels of the MCQs with their actual difficulty-levels is 79%.

Discussion and future work Even though the results of our difficulty-level predication method have shown a high correlation with the actual difficulty-levels, there are cases where the approach had failed to give a correct predication. In our observation, the repetition of similar words or part of a phrase in an MCQ's stem and its key, is one of the main reasons for this unexpected behavior. This word repetition can give a hint to the learner, enabling her to choose the correct answer. Example-4 shows the MCQ item i_{15}, where the repetition of the word "string" in the stem and in the key, and appearance of the stem component "Robin-Karp Algorithm" in the choice set has degraded the MCQ's (actual) difficulty-level. Since our approach is not designed to identify such flaws in the questions, the predicted characteristics of the MCQs may deviate from their actual characteristics.

Example 4. **Stem:** Choose a *String Matching Algorithm*, which *is faster than Robin-Karp Algorithm*.

Choice set: a. Robin-Karp algo. b. Bubble sort c. Boyer Moore string search algo.

A validity check based on the quality assurance guidelines of an MCQ question (suggested by Haladyna et al. in [18]) has to be done prior to finding the difficulty-levels of the MCQs. This would prevent the MCQs that have the the above mentioned flaws becoming part of the assessment test.

Furthermore, we observed that, out of 811 stems that are generated for the experiment, only 46% of the stems were found to be valid. An initial analysis of these invalid stems has shown that some of these stems are wrongly marked as invalid; showing *false negative* errors. A detailed analysis of these false negatives has to be done in future, to further refine the difficulty measures.

Another future work is based on our observation that, Eq. 1 for D_{expert} cannot correctly capture the difficulty-scores for all the cases. For example, when the $ASpace$ is very small, the equation tends to give a small difficulty-score, but (intuitively) in reality this may not be the case. When the question pattern becomes very rare then it will instead become difficulty for an expert to answer it correctly. Similar is the case when the $ASpace$ is too large, the equation gives a high difficulty-score, but in reality this may not be the case, as when the $ASpace$ becomes too large, the question becomes too generic that an expert will easily answer it. We came across similar observations for the equation which we used to calculate $D_{beginner}$. Due to the page limit, we refrain from giving more details. A detail empirical study is required to validate our observations.

8 Conclusion

We have proposed an item difficulty model for MCQs that are generated from formal ontologies. Various factors such as popularities of the ontology entities, answer space (number of instances in a class) and hierarchical depth of predicates, were considered as a part of the difficulty model for predicting MCQs' difficulty values. A criteria for

checking the validity of the predicated difficulty-scores is also detailed. A detailed empirical study under real world conditions (making use of the Item Response Theory principles) is conducted to find out the actual difficulty-levels of the automatically generated MCQs; a comparison of the predicated difficulty-levels with their actual difficulty-levels has shown a great correlation.

In this paper we have tried to characterize some of the factors that are associated with the difficulty of an MCQ question. Our experiments cannot conclusively show that the characterization of the factors is absolute or the factor which have been considered are the only factors which determine the difficulty-level of an MCQ. However, our empirical evaluation has shown that the measures which we proposed can correctly predict the difficulty-score of an ontology generated MCQ to a large extent.

Table 2: The probabilities of correctly answering the test MCQs (i.e., P values) and the α_i values calculated using the obtained P values are shown below. The learners in L_1, L_2 and L_3 are having high, medium and low domain knowledge proficiencies respectively.

MCQ item No.	P values for the learner sets			α_i values for the learner sets			Actual Diffi.-level	Predicted Diffi.-level
	L_1	L_2	L_3	L_1	L_2	L_3		
i_1	0.32	0.12	0.08	2.25	1.99	0.94	$high$	$high$
i_2	0.43	0.32	0.12	1.78	0.75	0.49	$high$	$high$
i_3	0.45	0.22	0.00	1.70	1.27	$+\infty$	$high$	$high$
i_4	0.51	0.42	0.06	1.46	0.32	1.25	$high$	$high$
i_5	0.58	0.27	0.00	1.18	0.99	$+\infty$	$high$	$high$
i_6	0.44	0.30	0.11	1.74	0.85	0.59	$high$	$high$
i_7	0.41	0.17	0.08	1.86	1.59	0.94	$high$	$high$
i_8	0.72	0.41	0.06	0.56	0.36	1.25	$--$	$high$
i_9	0.96	0.51	0.02	-1.68	-0.04	2.39	$medium$	$medium$
i_{10}	0.94	0.22	0.06	-1.25	1.24	1.25	$medium$	$medium$
i_{11}	1.00	0.46	0.00	$-\infty$	0.16	$+\infty$	$medium$	$medium$
i_{12}	0.90	0.56	0.11	-0.70	-0.24	0.59	$medium$	$medium$
i_{13}	1.00	0.61	0.08	$-\infty$	-0.45	0.94	$medium$	$medium$
i_{14}	1.00	0.44	0.00	$-\infty$	0.24	$+\infty$	$medium$	$medium$
i_{15}	1.00	0.89	0.40	$-\infty$	-2.09	-1.09	$--$	$medium$
i_{16}	1.00	0.48	0.06	$-\infty$	0.08	1.25	$medium$	$medium$
i_{17}	1.00	0.82	0.38	$-\infty$	-1.52	-0.84	$--$	low
i_{18}	1.00	0.84	0.16	$-\infty$	-1.66	0.16	$--$	low
i_{19}	1.00	1.00	0.48	$-\infty$	$-\infty$	-1.42	low	low
i_{20}	0.96	0.88	0.50	-1.68	-1.99	-1.50	low	low
i_{21}	1.00	0.92	0.32	$-\infty$	-2.44	-0.75	$--$	low
i_{22}	1.00	0.96	0.26	$-\infty$	-3.18	-1.05	low	low
i_{23}	1.00	0.77	0.22	$-\infty$	-1.21	-1.18	low	low
i_{24}	1.00	0.88	0.00	$-\infty$	-1.99	-1.26	low	low

References

1. Michael Furr, R., Bacharach, V.R.: Psychometrics, An Introduction. Second Edition. SAGE Publications, Inc (2014)

2. David Miller, M., Robert L., L., Gronlund, N.E.: Measurement and Assessment in Teaching, Tenth Edition. Pearson (2008)
3. Kehoe, J.: Basic Item Analysis for Multiple-Choice Tests. ERIC/AE Digest (ERIC Number: ED398237), ERIC Publications (1995)
4. Cubric, M., Tosic, M.: Towards automatic generation of e-assessment using semantic web technologies. In: Proceedings of the 2010 International Computer Assisted Assessment Conference. (2010)
5. Vinu, E.V., Kumar, P.S.: A novel approach to generate MCQs from domain ontology: Considering DL semantics and open-world assumption. Web Semantics: Science, Services and Agents on the World Wide Web **34** (2015) 40 – 54
6. Heilman, M., Smith, N.A.: Good question! statistical ranking for question generation. In: Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA. (2010) 609–617
7. Vinu, E.V., Kumar, P.S.: Improving large-scale assessment tests by ontology based approach. In: Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2015, Hollywood, Florida. May 18-20, 2015. (2015) 457
8. Alsubait, T., Parsia, B., Sattler, U.: A similarity-based theory of controlling mcq difficulty. In: e-Learning and e-Technologies in Education (ICEEE), 2013 Second International Conference on. (Sept 2013) 283–288
9. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F., eds.: The description logic handbook: theory, implementation, and applications. Cambridge University Press, New York, NY, USA (2003)
10. Papasalouros, A., Kanaris, K., Kotis, K.: Automatic Generation of Multiple Choice Questions from Domain Ontologies. In Nunes, M.B., McPherson, M., eds.: e-Learning, International Association for Development of the Information Society (2008) 427–434
11. Holohan, E., Melia, M., McMullen, D., Pahl, C.: Adaptive e-learning content generation based on semantic web technology. In: Proceedings of Workshop on Applications of Semantic Web Technologies for e-Learning, Amsterdam, The Netherlands. (2005) 29–36
12. Alsubait, T., Parsia, B., Sattler, U.: Mining ontologies for analogy questions: A similarity-based approach. Volume 849 of CEUR Workshop Proceedings., OWL Experiences and Directions (2012)
13. Alsubait, T., Parsia, B., Sattler, U.: Generating multiple choice questions from ontologies: Lessons learnt. In: Proceedings of the 11th International Workshop on OWL: Experiences and Directions (OWLED 2014). Volume 1265. (Oct 2014) 73–84
14. Alsubait, T., Parsia, B., Sattler, U.: Ontology-based multiple choice question generation. KI - Künstliche Intelligenz (2015) 1–6
15. Ben Abacha, A., Zweigenbaum, P.: Medical question answering: Translating medical questions into sparql queries. In: Proceedings of the 2Nd ACM SIGHIT International Health Informatics Symposium. IHI '12, New York, NY, USA, ACM (2012) 41–50
16. itko, B., Stankov, S., Rosi, M., Grubii, A.: Dynamic test generation over ontology-based knowledge representation in authoring shell. Expert Systems with Applications **36**(4) (2009) 8185 – 8196
17. Dehnad, A., Nasser, H., Hosseini, A.F.: A comparison between three-and four-option multiple choice questions. Procedia - Social and Behavioral Sciences **98** (2014) 398 – 403 Proceedings of the International Conference on Current Trends in ELT.
18. Haladyna, T.M., Downing, S.M., Rodriguez, M.C.: A review of multiple-choice item-writing guidelines for classroom assessment. Applied Measurement in Education **15**(3) (2002) 309–333+
19. Tartir, S., Arpinar, I.B., Moore, M., Sheth, A.P., Aleman-meza, B.: Ontoqa: Metric-based ontology quality analysis. In: IEEE Workshop on Knowledge Acquisition from Distributed, Autonomous, Semantically Heterogeneous Data and Knowledge Sources. (2005)

20. Revuri, S., Upadhyaya, S.R., Kumar, P.S.: Using domain ontologies for efficient information retrieval. In Lakshmanan, L.V.S., Roy, P., Tung, A.K.H., eds.: COMAD, Tata McGraw-Hill Publishing Company Limited (2006) 170–173